

Recursive Visual Programming

Jiaxin Ge¹ Sanjay Subramanian¹ Baifeng Shi¹
Roei Herzig¹ Trevor Darrell¹

UC Berkeley, CA, USA

Abstract. Visual Programming (VP) has emerged as a powerful framework for Visual Question Answering (VQA). By generating and executing bespoke code for each question, these methods show advancement in leveraging Large Language Models (LLMs) for complex problem-solving. Despite their potential, existing VP methods generate all code in a single function, which does not fully utilize LLM’s reasoning capacity and the modular adaptability of code. This results in code that is suboptimal in terms of both accuracy and interpretability. Inspired by human coding practices, we propose Recursive Visual Programming (RVP), which better harnesses the reasoning capacity of LLMs, provides modular code structure between code pieces, and assigns different return types for the sub-problems elegantly. RVP approaches VQA tasks with an top-down recursive code generation approach, allowing decomposition of complicated problems into smaller parts. We show RVP’s efficacy through extensive experiments on benchmarks including VSR, COVR, GQA, and NextQA, underscoring the value of adopting human-like recursive and modular programming techniques for solving VQA tasks. Our code is available at <https://github.com/para-lost/RVP>.

Keywords: Visual Programming · Vision-Language · VQA

1 Introduction

Visual Question Answering (VQA) lies at the intersection of computer vision and natural language processing, posing the challenge of interpreting visual data to answer questions [3, 16, 26]. Following the recent progress on code generation using Large Language Models (LLMs) [6, 27, 33], Visual Programming (VP) methods have notably advanced in this area, successfully using LLMs to generate and execute code in few-shot and zero-shot scenarios [14, 30, 31]. In these techniques, various vision models are provided to LLM as APIs, and the LLM is used as a planner to utilize these vision models to perform reasoning over images.

Despite these advancements, such planning remains very challenging for Visual Programming. These methods require the model to generate a single code block using the predefined APIs at once. As a result, the model must handle logic, be aware of all details, and utilize the vision APIs appropriately. This has two main drawbacks. First, such approaches do not fully utilize LLM’s reasoning capability, which would benefit from breaking down the problem into step-by-step sub-problems, as demonstrated in recent NLP research [4]. Secondly, they

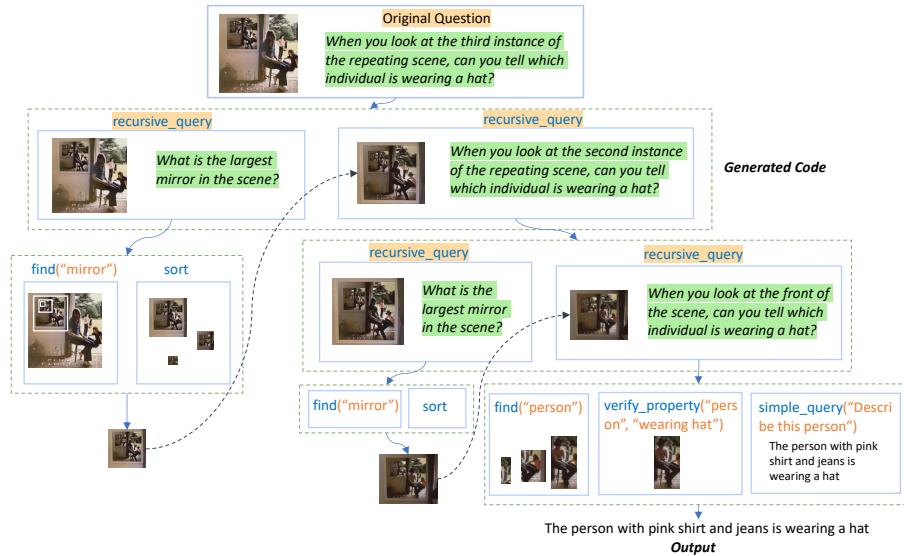


Fig. 1: A breakdown of Recursive Visual Programming for a visual question is illustrated with an image where scenes within it contain smaller versions of themselves. To locate a man wearing a hat at the third level of an image, the model first generates code that has two `recursive_query` API calls. The initial call identifies the largest mirror (representing the next image level), and the subsequent call seeks “the man with a hat” at the second level. The model iteratively generates and executes code through these calls until the final answer is produced without additional recursive calls.

fail to utilize the full potential of code because they lack modularity between the code pieces, particularly when the code is one long piece.

We consider how humans tackle these issues. Humans tend to code in a top-down manner: first they break down complex tasks into manageable sub-problems, and then they tackle each through focused code segments. This approach allows the programmer to concentrate on the main logic flow rather than figuring out all the details at once, resulting in a more clear and elegant structure.

Inspired by this, we hypothesize that incorporating a similar recursive coding strategy would also help solve complex visual question-answering problems. To investigate whether LLMs could indeed benefit from recursive coding, we conduct preliminary experiment of the recursive coding approach on standard symbolic reasoning benchmarks, such as Dyck Language [32] and Games of 24 [38]. This is done by simply prompting the LLM to code recursively. We find that utilizing LLMs to code recursively leads to a notable improvement in accuracy, as shown in Figure 2. Motivated by this result, we explore in this paper how a recursive coding methodology can be applied to VQA tasks.

Specifically, we propose a novel VP method for VQA tasks, named **Recursive Visual Programming (RVP)**. Instead of querying LLM only once to generate a single piece of code, we ask LLM to decompose the question into sub-questions




Fig. 2: Motivating Examples. Adding the prompt “solve the problem recursively” to the input significantly improves the model’s performance on traditional coding tasks, Dyck language [32] and Games of 24 [38]. “NR” refers to non-recursive approach and “R” refers to recursive approach. This shows that LLMs can potentially benefit from human’s recursive programming approach.

and recursively query itself for each sub-question to generate modular code pieces. For each sub-question, the same process repeats, until reaching an atom level where the current question can be solved by only calling external APIs. See example in Figure 1. In contrast to querying LLM only once, we find that using recursive queries to generate code for each sub-question can improve the readability and accuracy of the code, as shown in Figure 3. Moreover, RVP allows assigning dynamic types to each sub-question, enabling new code pieces to return a variety of types based on the current need. Remarkably, RVP assigns return types unseen in in-context examples, such as `List[str]`, `List[ImagePatch]`.

Our contribution can be summarized as follows: (i) We introduce a novel recursive approach based on human programming. This approach is a top-down programming method that better harnesses LLM’s reasoning capacity and enhances modularity between multiple code pieces. It is easily applicable to any existing VP model for visual question answering task. (ii) We allow dynamic return types for the sub-problems, which enhances flexibility and adaptability. (iii) We demonstrate improved performance on several standard benchmarks, such as GQA, VSR, NextQA and COVR, demonstrating the effectiveness of our approach. We also show better interpretability than non-recursive VP methods.

2 Related Work

Modular Visual Reasoning. There is a long line of work that seeks to combine modularity with the expressivity of deep neural networks. Early work on this topic converts the output of a parser [1] to a program or trains a program generator either separately from or jointly with the execution modules [2, 15]. Later work finds that such models are often challenging to train on real-world datasets, particularly when interpretability of individual module outputs is expected [29]. Other work improves performance on multi-hop visual reasoning without explicitly defining a set of modules, e.g., by predicting and operating

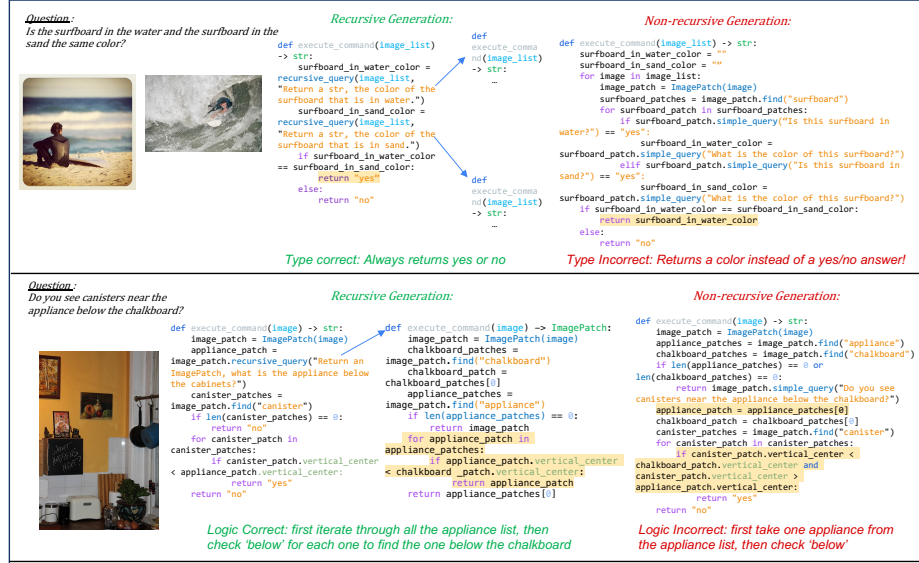


Fig. 3: Example from COVR (Upper). Current VP methods fail to provide binary 'yes' or 'no' answers, while Recursive Visual Programming method outputs the correct answer. **Example from GQA (Middle and Bottom).** Recursive outperforms current non-recursive methods by correctly addressing all details and their associated logic.

over a scene graph [18].

Visual Programming. Recent language models can accurately generate code from text descriptions when prompted with a small number of examples [8]. This advance has enabled a few-shot modular visual reasoning approach, commonly called visual programming [14], which consists of prompting a language model for a program and executing the program with pre-trained vision models. This paradigm has led to impressive results in visual question answering, video question answering, visual and tabular question answering, referring expression comprehension, and text-based image editing [11, 12, 23, 30, 31]. Programming with LLMs has also been applied to produce plans and reward functions for robots [21, 24, 40]. Our work builds on these techniques by enabling the LLM to generate code in multiple steps—first planning the overall process and then filling in the details—rather than all at once.

Language Models for Reasoning Tasks. In the text domain, there is a large body of work on using language models for reasoning tasks. In particular, several papers have shown the efficacy of code generation as a scaffold for such tasks [9, 13, 25, 34]. Aside from code generation, prior work has also shown that decomposing complex queries into simpler ones improves the performance of language models [28, 42]. Our RVP approach applies this insight to improve the code generation component of visual programming.

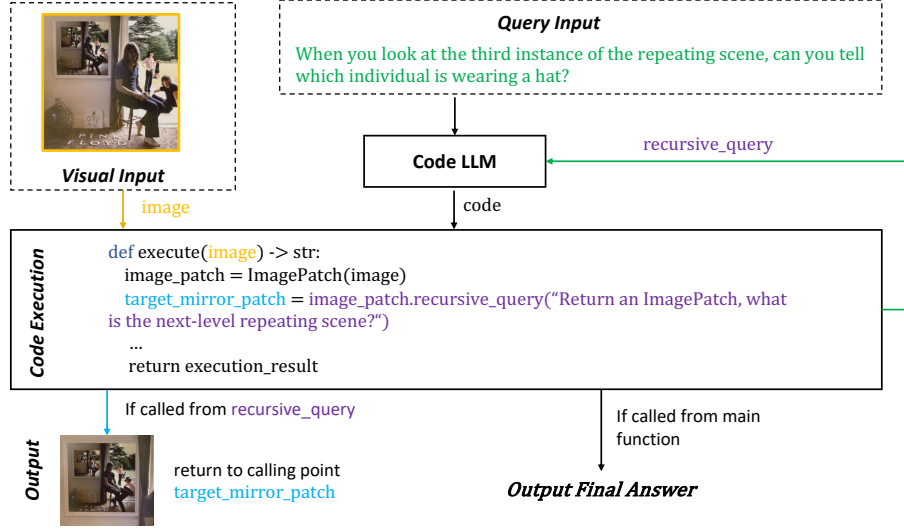


Fig. 4: An overview of our approach. When a `recursive_query` API call is occurred, the model generates and executes new code based on the query, returning the output back to the original call.

3 Recursive Visual Programming

In this section, we introduce our method, Recursive Visual Programming (RVP). We first review the problem setup and ingredients of visual programming as established by prior work (see Section 3.1). We then introduce our recursive coding framework (see Section 3.2), which involves a multi-step prompting scheme, as well as a system for dynamic type assignment (see Section 3.3).

3.1 Preliminaries

In Visual Question Answering (VQA), the system f is given a visual input I (*i.e.* an image, a set of images, or a video), as well as a question Q relating to the visual content. The expected output is a textual answer A to the question:

$$A \leftarrow f(I, Q) \quad (1)$$

Visual Programming methods [14, 30, 31] represent f as the composition of two functions: a code generator g and a visual executor h . In the first part, the code generator generates a program P conditioned on the question, and in the second part, the executor executes the program on the visual input:

$$P \leftarrow g(Q), A \leftarrow h(P, I) \quad (2)$$

We focus on the few-shot setting in which we have only a few (at most 50) training examples—each consisting of a question, program, and answer—for a given

task. Therefore, the code generator g is composed of a language model and a prompt that includes a description of the API available to the executor and in-context examples of questions along with the corresponding correct programs. We use Python as the programming language, similar to [30, 31].

We evaluate our approach using the ViperGPT [31] API. For questions over images or sets of images, these functions can be called on an `ImagePatch` object, which can be any crop of the image:

1. `find(object: str)`: Returns a list of `ImagePatch` objects matching the description.
2. `exists(object: str)`: Returns a bool representing whether the described object exists.
3. `verify_property(object: str, property: str)`: Returns a bool representing whether the described object has the specified property.
4. `simple_query(question: str)`: Returns an answer to the provided question based on the given `ImagePatch`.
5. `compute_depth()`: Returns the median depth of the region in the given `ImagePatch`.
6. `crop(left: int, lower: int, right: int, upper: int)`: Crops a part of the given `ImagePatch` and returns a new `ImagePatch` representing that crop.

For questions about a video, the API includes three other functions: `trim()`, `frame_from_index()`, and `frame_iterator()`, which return an abbreviated version of the video, a specific frame, and an iterator over the frames respectively.

3.2 Recursive Visual Programming

We introduce a new method for code generation in visual programming: Recursive Visual Programming (RVP). RVP enables the language model to break the code generation process into multiple sub-questions and recursively call itself to generate a program for each sub-question. As shown in Figure 4, we operationalize RVP by adding a new function to the API: `recursive_query(image, sub-question)`. During execution, when `recursive_query` is called, it prompts the language model to generate a new code given the sub-question. This code is then executed, and its return value is returned by `recursive_query`. An example break down of the usage of `recursive_query` call is shown in Figure 1. We set the recursive termination condition by (1) setting max depth of 10 and (2) switching to direct simple query when generated code contains a recursive call to the same input question.

To enable this process, we add examples of calling `recursive_query` in the prompt to show the model how to decompose the problem and utilize the answers returned by `recursive_query`. We provide examples of the in-context prompt of `recursive_query` we used in the supplementary material.

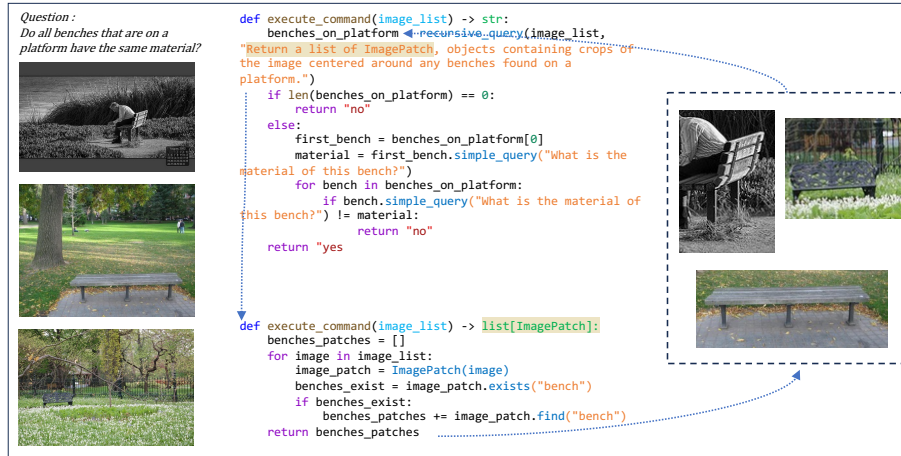


Fig. 5: Examples of dynamic type assignment in recursive_query calls. The model autonomously determines the appropriate return type in its code generation and generalize to new types unseen in in-context examples such as `List[ImagePatch]`.

3.3 Dynamic Type Assignment

In previous visual programming frameworks, code generation mechanisms are constrained by static type assignments, *i.e.*, the return type of a given function is fixed. For instance, `simple_query` always returns a string. However, this rigid structure limits the model’s flexibility to adapt the code to different contexts within the VQA task. RVP introduces dynamic type assignment to address this limitation. When the code generator writes a `recursive_query` function, it can include the expected return type as a part of the question that is passed to the `recursive_query`. To facilitate this, in the in-context examples that include `recursive_query`, we let the input ‘question’ specify a type. As shown in Figure 5, we give examples of type `ImagePatch`, `bool`, `float`, and `str` being specified within the queries. When receiving this question, the model discerns the required type and generates code that declares and returns the correct type. For instance, a query like “Return a bool, is there a black cat in the picture?” will prompt the generation of code with the signature.

$$\text{execute}(\text{image}) \rightarrow \text{bool} \quad (3)$$

This signature indicates that the generated code P' , when executed, will return a boolean value A_{bool} which answers the sub-question Q' , illustrating the dynamic and context-sensitive nature of RVP’s code generation capabilities that help to produce syntactically and semantically appropriate code for complex reasoning.

4 Experiments

In this section, we introduce the basic settings including the task formulation and parameter details. Then, we provide experiment results and in-depth analysis.

4.1 Experiment Setup

Datasets. We evaluate on four datasets spanning a wide range of visual reasoning skills and types of visual input: Visual Spatial Reasoning (VSR) [22], GQA [17], COVR [5], and NextQA [36]. For NextQA, we present results on the “Hard Temporal” subset which was curated to consist of questions that truly require temporal reasoning [7]. VSR and GQA include queries about single images, while COVR includes queries about multiple images, and NextQA includes queries about video. On each dataset, we evaluate methods by their accuracy in producing an answer that exactly matches the ground-truth answer. For NextQA, we adopt a multiple-choice setup in which we provide methods with the set of possible answer choices for each question. For VSR, a spatial statement of an image is made, and the question is framed as : “Is it true or false,” followed by the statement. In this way, a VSR statement is framed as a VQA question and the model will answer either true or false.

Implementation Details. We implement RVP using the ViperGPT API and the accompanying prompt template. For each dataset, we curate a set of in-context examples which are added to the prompt. The set of question-answer pairs is shared between the recursive and non-recursive methods. We use seven in-context examples for GQA, six for NextQA, six for VSR, and eight for COVR. For code generation, use GPT-3.5-turbo via the OpenAI API ¹.

4.2 Quantitative Results

Table 1 presents our results across four datasets. In all settings, we outperform ViperGPT, which is identical to RVP except that RVP includes `recursive_query` function in the API and in-context examples. This indicates that recursive programming indeed improves accuracy in visual reasoning tasks. Moreover, our method outperforms previous zero/few-shot methods in all but one of the settings. While our primary aim was not to push the state-of-the-art results, we show that integrating the clarity and elegance of recursive coding into standard VP methods not only doesn’t impair performance, but can enhance it.

4.3 Qualitative Results

We provide qualitative examples in Figure 3,6,7. In the example from COVR in Figure 3, the non-recursive approach answers with neither “yes” nor “no,” while the recursive approach does confirm to these choices. One possible explanation for this issue is that LLMs struggle to maintain coherence across long contexts [37]. RVP circumvents this issue by pushing the LLM to commit to the answer

¹ www.platform.openai.com

Table 1: Results. We report exact-match accuracy on each dataset. Best zero/few-shot result for each dataset is **bolded**. †Fully supervised results are from different models: ViLT [19] for VSR, HiTeA [39] for NextQA, VinVL-Base [41] for GQA, and VisualBERT [20] for COVR. ‡We report the results of our reproduction of ViperGPT using the officially released code and the same in-context examples used for our method (since the original in-context examples from ViperGPT were not released). We use GPT3.5-turbo while CodeVQA uses Codex, which is not available anymore.

Method	VSR		NextQA		GQA	COVR
	Random Split	Zero-shot Split	Hard Split-T	Test-dev	Test-dev	Test
Fully supervised†	69.3	63.0	48.6	65.1	57.9	
<i>Zero/few-shot methods</i>						
CLIP	56.0	54.5	—	—	—	—
BLIPv2	—	—	—	42.31	—	—
CodeVQA	—	—	—	49.0	50.7	—
ViperGPT‡	61.25	61.59	47.21	44.63	51.69	—
RVP (ours)	63.53	66.09	48.82	45.62	52.67	—

types earlier in its generation (the length from the beginning of the code to the point “return xxx” is shorter). In Figure 6, the non-recursive approach fails to handle the logic correctly while the recursive approach correctly does. This is likely due to that RVP breaks down the code logic and each code piece only needs to focus on a single logic, while the non-recursive method attempts to handle all the logic at once. In the GQA example, a common pattern involves handling complex logical problems and multiple properties. Non-recursive methods often fell short in these cases, either overlooking critical details or processing them in a logically flawed manner. RVP manages these intricate questions by addressing each detail accurately. This difference may stem from the inherent difficulty LLMs face in processing complicated logic over extended contexts. Recursive methods simplify this by focusing on the overall logic and gradually breaking the problems down into smaller, more manageable segments. This facilitates easier and more accurate resolution.

We provide a dynamic type assignment example in Figure 5. In this example, RVP assigns the return type `List[ImagePatch]` in the `recursive_query` call. It’s worth noting that the return type `List[ImagePatch]` is unseen in the in-context examples. This suggests that RVP has the potential to utilize complicated and diverse data structures, highlight the effectiveness of our approach.

4.4 Ablation Study

Dynamic Type This section studies integrating dynamic type assignment within the `recursive_query` API. For the GQA dataset, we examined all instances with recursive patterns in the `test_dev` split. In COVR, where the test set is unavailable, we used the `val_1000` samples from the validation set, following [30].

We compare four approaches: **Non-recursive Baseline**: Standard approach without recursive features. **Fixed-Type Recursive**: The `recursive_query` API is constrained to return a string. **Dynamic-Implicit Recursive**: The API

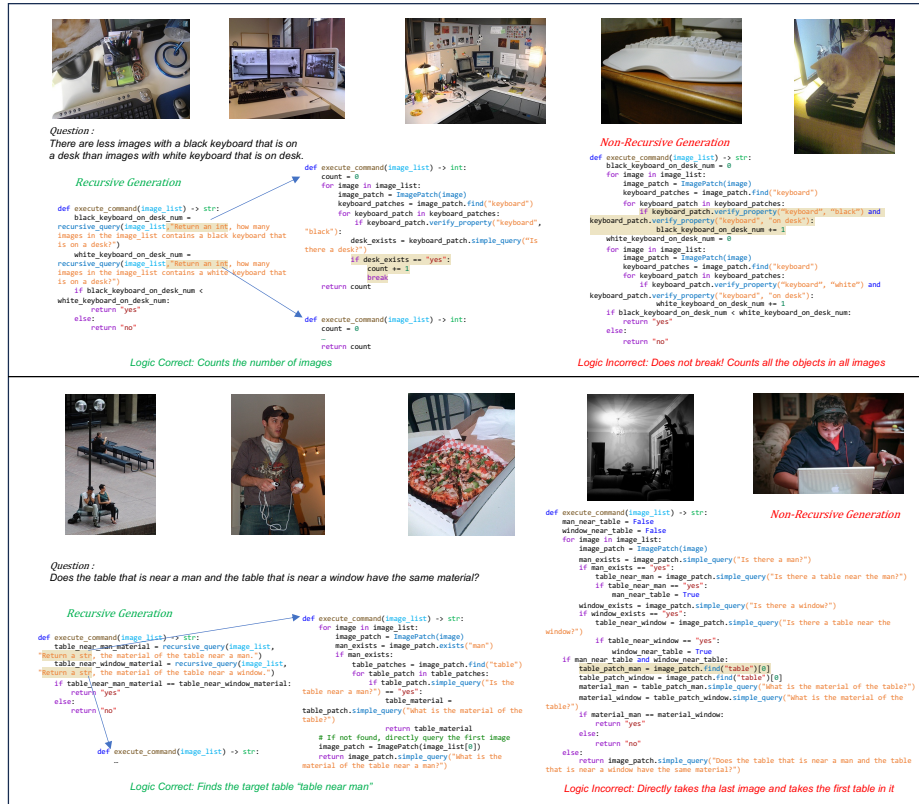


Fig. 6: Examples from COVR. RVP allows correct logic flow while traditional non-recursive VP fails to handle the logic correctly.

can return various types, but the specific type is not predefined in the query. For example, "Is there a black cat?" could yield either a boolean or a string. **Dynamic-Explicit Recursive:** Our primary approach, where the API's input query specifies the return type, like "Return a bool, is there a black cat?".

Table 2 shows our findings. The dynamic-explicit type assignment outperformed other methods, underscoring the need for clear type specification. The dynamic-implicit method showed lower performance than non-recursive due to type ambiguities, as exemplified in Figure 8, where the first code piece assumes the return type is a bool, while the second code piece returns a str. Fixed-type assignment proves to be beneficial over non-recursive methods, indicating that recursion itself adds value, but its rigidity in type handling limits performance. This study highlights the benefits of flexible, dynamic type assignment and the crucial role of explicit type declaration for accurate and consistent responses. These insights mark a step towards more adaptable visual reasoning frameworks, emphasizing the need for precision in dynamic environments.



Fig. 7: Examples from GQA. RVP allows handling the details more elegantly and produces logically correct code.

In-Context Examples We conduct ablation on using retrieval-based in-context example selection and using different number of in-context examples.

Retrieval-based in-context example selection. We manage to write 9 programs using the recursive method from the 50 provided by CodeVQA [30], and add our original 3 recursive programs to them. We keep the non-recursive in-context example unchanged and choose 3 recursive in-context examples from these 12 for each question. We follow [35] to use the embedding-based retrieval technique that calculates a score for each example, and then select the top 3 examples for each question. The results of GQA_val_2000 are in Table 2. We find that the recursive pattern occurrence rate is significantly higher than using fixed in-context examples. This is likely because by retrieving similar recursively-decomposed questions, the model can more easily simulate their pat-

terns, thereby attempting to solve more problems recursively. However, the overall accuracy does not improve. This suggests that although the retrieval-based method encourages more diverse and suitable patterns of recursive decomposition, there is a need for improved techniques in selecting the example pool—such as focusing on those most suitable for recursive decomposition instead of randomly selecting—along with careful tuning of the example programs.

Number of in-context examples. We randomly add one example/delete one example from the `covr_1000_val`, the results are shown in Table 2. The results indicate that adding the number of examples can enhance the performance, but does not have a significant impact.

Error feedback loop Recent studies have studied the debugging ability of LLMs on code [10]. We study whether modular visual code structure enables better bug identification for visual programming. We sample 20 incorrect programs for 10 questions that have type/logic errors, and use GPT3.5 to (1) identify the bug, and (2) write the correct code given the bug. For identification, the model correctly finds 7 bugs correctly for RVP and 2 for non-recursive visual programming. For correction, the model fixes 7 bugs correctly for RVP and 7 for visual programming. This suggests that clearer and more modular code structure, as presented RVP, can potentially aids bug identification.

Table 2: Ablation study. (a) We provide ablation study results for dynamic type assignment and found that specified dynamic type assignment in RVP achieves the highest performance. (b) We provide results on using retrieval-based in-context examples and fixed examples. Retrieval-based method enables more diverse recursive patterns but requires careful choosing the prompt pool. (c) We provide results of using different number of examples and found this doesn’t influence performance significantly.

(a) Dynamic Type Ablation.			(b) Retrieval-Based and Fixed Examples		(c) In-Context Example Num
Method	GQA	COVR	Fixed Retrieval		Method Accuracy
Non-Recursive	55.99	60.71	Accuracy	54.75 53.79	Del 1 50.65
Fixed-Type	65.05	60.12	Rec Rate	3.55% 14.85%	Original 51.05
Dynamic-Implicit	55.34	53.57	Rec Acc	71.83 57.91	Add 1 51.35
Dynamic-Explicit	70.23	67.86			

4.5 Readability Study

This study investigates how recursive coding impacts code readability and comprehension for human programmers. Our goal is to move from complex code structures to more modular, clear, and understandable code. To evaluate this, we focus on code quality in terms of its understandability.

From the COVR dataset, we randomly select one example from each five different question types where recursive patterns occur. We present the main function of both their recursive and non-recursive code versions to 15 participants and ask them to grasp the high-level intent of the code as quickly as

```
def execute(image) -> str:
    image_patch = ImagePatch(image)
    is_black_skateboard =
    image_patch.recursive_query("Are there any black
    skateboards?")
    is_motorcycle = image_patch.recursive_query("Are
    there any motorcycles?")
    return bool_to_yesno(is_black_skateboard or
    is_motorcycle)

def execute(image) -> str:
    image_patch = ImagePatch(image)
    skateboard_patches = image_patch.find("skateboard")
    for skateboard_patch in skateboard_patches:
        if
        skateboard_patch.verify_property("skateboard",
        "black"):
            return "yes"
    return "no"
```

Return Type Confusion: Expect Bool but Actually Str.

Fig. 8: Failure example of dynamic type assignment due to unspecified type. The model confuses the type and return a str instead of a bool.




Fig. 9: Comparison of average understanding time for recursive and non-recursive codes. RVP demonstrates better readability and is easy to understand.

possible. For the non-recursive case, the main function is just the single code snippet generated by the LLM. For the recursive case, the main function consists of the code generated by the first LLM call (the high level code) but not the functions generated by the recursive queries. This helps us measure how fast users can understand the code's intention by looking at the main function, and reflects the self-explanatory nature of the code. Figure 9 shows the results. Participants consistently understand the recursive code faster than the non-recursive versions. This quicker comprehension can be attributed to the recursive code's structured and clear approach. Non-recursive code, while correct, sometimes lacks immediate clarity in conveying its primary goal.

We also calculate the average line length for all main and sub functions. For RVP, the average length of all functions is **8.47**. For non-recursive VP, the average length of the functions is **15.4**. This suggests that RVP decreases the average length the code pieces and enhances modularity.

Our findings suggest that recursive coding not only simplifies complex code but also enhances its accessibility and ease of understanding. Its orderly structure and clear logic accelerate comprehension, leading to more efficient and user-friendly programming. We provide more details of this study in the supp.

4.6 Recursive Analysis

Question Type. We study which types of questions tend to use recursive and found that recursive patterns appeared in five question types within both the COVR test set and the GQA test-dev set, as shown in Table 3. The 'Verify Property' and 'Query' types in GQA, and 'Compare' and 'Multi-Reference' in COVR, were most common. This suggests recursive methods tend to occur in tasks that require comparative and referential analysis.

Dynamic Type. We study which types of dynamic return types occurs and

Table 3: Recursive analysis. Table (a) provides the distribution of recursive patterns by question type in GQA and COVR. We have found that recursive patterns occur commonly in question types that require reasoning. Table (b) studies different recursive query return types and their frequency in GQA and COVR. The model generalizes to unseen types. Table (c) compares accuracy across question types. Recursive methods outperform non-recursive particularly in ‘Choose’ and ‘Query’ questions.

(a) Dist. of Recursive Patterns (b) Diverse Return Types (c) Type Accuracy

COVR/GQA Question Types			Data Type GQA COVR			Question Accuracy (%)		
	COVR	GQA					Non-Rec	Rec
Compare	457	15	Str	319	1828	Verify	52.94	64.71
Mult. Ref	282	-	Bool	521	592	Query	56.89	71.86
Quant. Attr	12	-	List[str]	16	-	Logical	55.88	70.59
Spec Attr	357	-	ImagePatch	30	-	Compare	66.67	53.33
Choose	1	25	List[ImagePatch]	1	5	Choose	52.00	84.00
Verify	-	68	Float	6	-			
Query	-	167	Int	-	914			
Logical	-	34						

their frequently in GQA and COVR. In Table 3, we found that RVP generalizes to new types like `List[str]` and `List[ImagePatch]`, which are not included in the in-context examples. This flexibility and ability to generalize beyond in-context examples highlight the inherent adaptability of recursive methods.

Type Accuracy. We evaluate the accuracy of recursive methods across different question types. An analysis of the GQA test-dev set is presented in Table 3. It reveals that recursive methods significantly improve accuracy in ‘Query’, ‘Logical’, and ‘Choose’ questions. This underscores their suitability for complex reasoning and decision-making tasks.

Error Source. We sample 50 incorrect recursive examples from `GQA2000val` and manually check the error source: 22% are caused by logic error inside sub-functions; 18% are caused by logic error between functions; 52% are caused by API model inaccuracy (i.e. detection model error); and 4% are caused by other problems like the groundtruth answer unclear.

Cost And Runtime. For the questions which require recursive query, RVP takes about 2.1 times (runtime) and about 3 times (cost) compared to ViperGPT.

5 Conclusion and Limitation

In this work, we present RVP, a visual programming method that adopts recursive coding to generate concise, elegant and easy-to-understand code, handling details and intricate logic more accurately compared with traditional visual programming methods. We conduct extensive experiments across various benchmarks and provide a comprehensive analysis of RVP. Our work is the first to explore recursive programming in the visual domain and dynamic-type code generation for visual reasoning; we believe the recursive programming concept would have a potential greater impact on other domains as well. While RVP is effective for complicated visual questions, most questions in current VQA benchmarks tend to be short and non-recursive. We believe that the benefit of our method could be demonstrated better with a new benchmark or a more complex split.

Acknowledgements

We would like to thank Ben Bogin for helping us test on the private COVR test set, and Guangyuan Jiang for helping us proof read the paper and providing valuable feedback.

References

1. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Neural module networks. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 39–48 (2015), <https://api.semanticscholar.org/CorpusID:5276660>
2. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Learning to compose neural networks for question answering. ArXiv **abs/1601.01705** (2016), <https://api.semanticscholar.org/CorpusID:3130692>
3. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: Vqa: Visual question answering. In: Proceedings of the IEEE international conference on computer vision. pp. 2425–2433 (2015)
4. Besta, M., Blach, N., Kubicek, A., Gerstenberger, R., Podstawski, M., Gianinazzi, L., Gajda, J., Lehmann, T., Niewiadomski, H., Nyczyk, P., Hoefler, T.: Graph of thoughts: Solving elaborate problems with large language models (2024)
5. Bogin, B., Gupta, S., Gardner, M., Berant, J.: Covr: A test-bed for visually grounded compositional generalization with real images. ArXiv **abs/2109.10613** (2021), <https://api.semanticscholar.org/CorpusID:237592834>
6. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 1877–1901. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>
7. Buch, S., Eyzaguirre, C., Gaidon, A., Wu, J., Fei-Fei, L., Niebles, J.C.: Revisiting the “Video” in Video-Language Understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
8. Chen, M., Tworek, J., Jun, H., Yuan, Q., Ponde, H., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F.P., Cummings, D.W., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W.H., Nichol, A., Babuschkin, I., Balaji, S.A., Jain, S., Carr, A., Leike, J., Achiam, J., Misra, V., Morikawa, E., Radford, A., Knight, M.M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., Zaremba, W.: Evaluating large language models trained on code. ArXiv **abs/2107.03374** (2021), <https://api.semanticscholar.org/CorpusID:235755472>
9. Chen, W., Ma, X., Wang, X., Cohen, W.W.: Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Transactions on Machine Learning Research* (2023)

10. Chen, X., Lin, M., Schärli, N., Zhou, D.: Teaching large language models to self-debug. arXiv preprint arXiv:2304.05128 (2023)
11. Cheng, Z., Xie, T., Shi, P., Li, C., Nadkarni, R., Hu, Y., Xiong, C., Radev, D.R., Ostendorf, M., Zettlemoyer, L., Smith, N.A., Yu, T.: Binding language models in symbolic languages. ArXiv **abs/2210.02875** (2022), <https://api.semanticscholar.org/CorpusID:252734772>
12. Cho, J., Zala, A., Bansal, M.: Visual programming for text-to-image generation and evaluation. NeurIPS (2023)
13. Gao, L., Madaan, A., Zhou, S., Alon, U., Liu, P., Yang, Y., Callan, J., Neubig, G.: Pal: Program-aided language models. arXiv preprint arXiv:2211.10435 (2022)
14. Gupta, T., Kembhavi, A.: Visual programming: Compositional visual reasoning without training. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 14953–14962 (2022), <https://api.semanticscholar.org/CorpusID:253734854>
15. Hu, R., Andreas, J., Rohrbach, M., Darrell, T., Saenko, K.: Learning to reason: End-to-end module networks for visual question answering. 2017 IEEE International Conference on Computer Vision (ICCV) pp. 804–813 (2017), <https://api.semanticscholar.org/CorpusID:18682>
16. Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 6693–6702 (2019), <https://api.semanticscholar.org/CorpusID:152282269>
17. Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 6693–6702 (2019), <https://api.semanticscholar.org/CorpusID:152282269>
18. Hudson, D.A., Manning, C.D.: Learning by abstraction: The neural state machine. In: Neural Information Processing Systems (2019), <https://api.semanticscholar.org/CorpusID:195847902>
19. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. In: International Conference on Machine Learning (2021), <https://api.semanticscholar.org/CorpusID:231839613>
20. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: Visualbert: A simple and performant baseline for vision and language. ArXiv **abs/1908.03557** (2019), <https://api.semanticscholar.org/CorpusID:199528533>
21. Liang, J., Huang, W., Xia, F., Xu, P., Hausman, K., Ichter, B., Florence, P., Zeng, A.: Code as policies: Language model programs for embodied control. In: arXiv preprint arXiv:2209.07753 (2022)
22. Liu, F., Emerson, G.E.T., Collier, N.: Visual spatial reasoning. Transactions of the Association for Computational Linguistics **11**, 635–651 (2022), <https://api.semanticscholar.org/CorpusID:248496506>
23. Lu, P., Peng, B., Cheng, H., Galley, M., Chang, K.W., Wu, Y.N., Zhu, S.C., Gao, J.: Chameleon: Plug-and-play compositional reasoning with large language models. arXiv preprint arXiv:2304.09842 (2023)
24. Ma, Y.J., Liang, W., Wang, G., Huang, D.A., Bastani, O., Jayaraman, D., Zhu, Y., Fan, L., Anandkumar, A.: Eureka: Human-level reward design via coding large language models. arXiv preprint arXiv: Arxiv-2310.12931 (2023)
25. Madaan, A., Zhou, S., Alon, U., Yang, Y., Neubig, G.: Language models of code are few-shot commonsense learners. In: Goldberg, Y., Kozareva, Z., Zhang, Y. (eds.) Proceedings of the 2022 Conference on Empirical Methods in Natural Language

- Processing. pp. 1384–1403. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Dec 2022). <https://doi.org/10.18653/v1/2022.emnlp-main.90>, <https://aclanthology.org/2022.emnlp-main.90>
26. Marino, K., Rastegari, M., Farhadi, A., Mottaghi, R.: Ok-vqa: A visual question answering benchmark requiring external knowledge. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 3190–3199 (2019), <https://api.semanticscholar.org/CorpusID:173991173>
 27. OpenAI: Gpt-4 technical report. ArXiv **abs/2303.08774** (2023), <https://api.semanticscholar.org/CorpusID:257532815>
 28. Press, O., Zhang, M., Min, S., Schmidt, L., Smith, N.A., Lewis, M.: Measuring and narrowing the compositionality gap in language models. ArXiv **abs/2210.03350** (2022), <https://api.semanticscholar.org/CorpusID:252762102>
 29. Subramanian, S., Bogin, B., Gupta, N., Wolfson, T., Singh, S., Berant, J., Gardner, M.: Obtaining faithful interpretations from compositional neural networks. In: Annual Meeting of the Association for Computational Linguistics (2020), <https://api.semanticscholar.org/CorpusID:218487535>
 30. Subramanian, S., Narasimhan, M., Khangaonkar, K., Yang, K., Nagrani, A., Schmid, C., Zeng, A., Darrell, T., Klein, D.: Modular visual question answering via code generation. In: Rogers, A., Boyd-Graber, J., Okazaki, N. (eds.) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 747–761. Association for Computational Linguistics, Toronto, Canada (Jul 2023). <https://doi.org/10.18653/v1/2023.acl-short.65>, <https://aclanthology.org/2023.acl-short.65>
 31. Sur’is, D., Menon, S., Vondrick, C.: Vipergpt: Visual inference via python execution for reasoning. ArXiv **abs/2303.08128** (2023), <https://api.semanticscholar.org/CorpusID:257505358>
 32. Suzgun, M., Scales, N., Schärli, N., Gehrmann, S., Tay, Y., Chung, H.W., Chowdhery, A., Le, Q.V., Chi, E.H., Zhou, D., et al.: Challenging big-bench tasks and whether chain-of-thought can solve them. arXiv preprint arXiv:2210.09261 (2022)
 33. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., Lample, G.: Llama: Open and efficient foundation language models. ArXiv **abs/2302.13971** (2023), <https://api.semanticscholar.org/CorpusID:257219404>
 34. Wang, X., Li, S., Ji, H.: Code4struct: Code generation for few-shot structured prediction from natural language. arXiv preprint arXiv:2210.12810 (2022)
 35. Wang, Z., Li, M., Xu, R., Zhou, L., Lei, J., Lin, X., Wang, S., Yang, Z., Zhu, C., Hoiem, D., et al.: Language models with image descriptors are strong few-shot video-language learners. Advances in Neural Information Processing Systems **35**, 8483–8497 (2022)
 36. Xiao, J., Shang, X., Yao, A., Chua, T.S.: Next-qa: Next phase of question-answering to explaining temporal actions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9777–9786 (June 2021)
 37. Yang, K., Klein, D., Peng, N., Tian, Y.: Doc: Improving long story coherence with detailed outline control. In: Annual Meeting of the Association for Computational Linguistics (2023), <https://api.semanticscholar.org/CorpusID:254877751>
 38. Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T.L., Cao, Y., Narasimhan, K.: Tree of thoughts: Deliberate problem solving with large language models. arXiv preprint arXiv:2305.10601 (2023)

39. Ye, Q., Xu, G., Yan, M., Xu, H., Qian, Q., Zhang, J., Huang, F.: Hitea: Hierarchical temporal-aware video-language pre-training. ArXiv **abs/2212.14546** (2022), <https://api.semanticscholar.org/CorpusID:255340506>
40. Yu, W., Gileadi, N., Fu, C., Kirmani, S., Lee, K.H., Gonzalez Arenas, M., Lewis Chiang, H.T., Erez, T., Hasenclever, L., Humplik, J., Ichter, B., Xiao, T., Xu, P., Zeng, A., Zhang, T., Heess, N., Sadigh, D., Tan, J., Tassa, Y., Xia, F.: Language to rewards for robotic skill synthesis. Arxiv preprint arXiv:2306.08647 (2023)
41. Zhang, P., Li, X., Hu, X., Yang, J., Zhang, L., Wang, L., Choi, Y., Gao, J.: Vinvl: Revisiting visual representations in vision-language models. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 5575–5584 (2021), <https://api.semanticscholar.org/CorpusID:235692795>
42. Zhou, D., Scharli, N., Hou, L., Wei, J., Scales, N., Wang, X., Schuurmans, D., Bousquet, O., Le, Q., hsin Chi, E.H.: Least-to-most prompting enables complex reasoning in large language models. ArXiv **abs/2205.10625** (2022), <https://api.semanticscholar.org/CorpusID:248986239>